In the Trenches with IP Multicast

Greg Bell
LBLnet Services Group, LBNL
ESCC: February 6, 2003

Goals for this talk

- To tell the story of our “multicast meltdown”
- To describe a few insidious failure modes
- To summarize lessons learned
- To offer tips for faster troubleshooting
- Caveats:
  - This is not a traditional guide to debugging multicast
  - I’m describing a worst-case scenario
  - I have an *operational* perspective

A brief orientation

- **LBLnet is a medium-sized LAN**
  - **>12,000 attached devices**
  - **>100 subnets**
  - **IP, IPX, Appletalk & DECnet**
  - **several remote sites**
- **Simple multicast topology**
  - **PIM sparse-dense, with MSDP for peering**
- **Major customer for multicast services: Access Grid Node**

Timeline

- until 9/01: stability
- 9/01-7/02: sporadic trouble, with 3 internal and 4 external outages
- 8/02-9/02: meltdown
  - Frequent multicast outages
  - Symptoms: loss of PIM neighbor state, severe packet loss (both random and cyclic), unidirectional connectivity, late joins, MSDP anomalies, and more…
  - Furious debugging effort
  - Time required to sort it out = engineer-weeks
  - Eventually diagnosed at least 8 bugs on 5 platforms

The happy ending

- **Return of stability on 10/7/02, when we diagnosed an intermittent CGMP bug**
- **No multicast failure within LBLnet since that date**
  - **(with one exception: the "SQL Slammer" worm tickled a Cisco MSDP bug, but the work-around was simple)**

Contributing to the meltdown

- Heterogeneous network (4 router platforms)
- Many IOS versions (the result of bug fixes)
- During the Summer of 2002, there was a significant increase in
  - AG Node utilization
  - The profile and importance of AG Node meetings
  - Volume of multicast traffic
- New peering relationship with ESnet
- New border topology
- All of these factors put us at greater risk

Lessons learned

- **Lesson #1: multicast routing code is buggy**
  - **We discovered catastrophic bugs on 5 different platforms (a mixture of routers / switches):**
    - **Cisco 4500**
    - **Cisco 7513**
    - **Cisco 6509**
    - **Cisco 8540**
    - **Cisco 5500**

How serious are the bugs?

- **Extremely serious (not minor annoyances). Often, they affect router stability, not just multicast functionality**
- **From a major vendor's website:**

> **Bug ID #CSCdx82485**
>
> **Symptoms: Under rare circumstances, a router that is configured with Protocol-Independent Multicast (PIM) may pause indefinitely.**
>
> **Workaround: Use a different Ethernet card, or avoid using PIM.**

How serious are the bugs (cont)?

- **We've seen all of the following symptoms:**
  - **Router reboots when it encounters multicast traffic**
  - **Normal PIM hello packets wedge an interface buffer, causing router to lose all PIM neighbor state**
  - **Router becomes unstable and drops unicast traffic when normal IGMP packets (mtrace) wedge a backbone interface buffer**
  - **Router spontaneously reboots when attempting to establish MSDP peering**
  - **MSDP router doesn't advertise active local sources**
  - **Switch running CGMP intermittently drops all multicast packets**

Lessons learned

- **Lesson #2: the severity of these bugs suggests a flagrant lack of concern for quality assurance**
  - **If PIM neighbor state isn't verified, then what is?**
  - **Beware! OS versions recommended by a vendor for multicast bug-fix purposes may contain new (and even worse) bugs**

Lessons learned

- **Lesson #3: stability *can* be achieved, but it may require considerable engineering resources**
- **Lesson #4: debugging a serious multicast problem may impair the stability of unicast routing**
  - **frequent OS upgrades**
  - **CPU-intensive debug commands**
  - **intrusive tests (process-switching of multicast)**

Lessons learned

- **Lesson #5: your "problem" may be caused by 2 (or more) simultaneous bugs, in which case troubleshooting becomes much more difficult**
  - **One symptom masks another**
  - **It's not always obvious when you've eliminated a bug**
- **Lesson #6: your diagnostic tools may be flawed**
  - **mtrace destabilizes remote router**
  - **Some "show" commands have bugs (ie, sho ip mroute count)**

Lessons learned, cont.

- **Lesson #7: when multicast fails, it will fail at the worst possible moment**
  - **during a meeting in which network budgets are discussed**
  - **during a high-profile Earth Simulator conference**
  - **2 minutes before the Lab Director arrives at the AG Node for a demonstration**
- **Not simply the result of bad luck; many bugs are load-dependent**
- **Lesson #7 leads to a special state of mind: Multicast Induced Paranoia (MIP)**

Practical suggestions

- **Once multicast is stable, what can you do to accelerate the trouble-shooting process?**

The key is rapid diagnosis

- **Easier said than done, because symptoms are frequently misleading**
  - **Problems in the WAN can initially present as problems in the LAN, and vice-versa**
  - **Once your network has suffered from multicast trouble, all AG Node problems are likely to be blamed on multicast**
    - **OS and application bugs**
    - **Sound and video-card issues**
    - **Cabling problems**
    - **Operator error**
- **Holy grail = rapid fault isolation. But how to achieve that?**

Practical suggestions

- **Consider putting major multicast applications (AG Node, for example) on a dedicated subnet**
  - **A case of theoretical purity vs. operational urgency**
  - **Less likely that debugging multicast will adversely affect other users**
- **Eliminate IGMP-snooping / CGMP on this subnet**
  - **Fewer failure modes**

Practical suggestions, cont.

- **Install a completely redundant host on the AG Node subnet to help rule out system-related trouble**
  - **Laptop running PIG or inSORS software**
- **Create a clear troubleshooting protocol for AG Node operators**
  - **When & how do they hand off a problem?**
- **Install copper or fiber taps for monitoring multicast traffic**
  - **“show” commands may deceive**
  - **tcpdump is your friend**
  - **Easy to do in conjunction with Bro**

Practical suggestions

- **Use the multicast beacon, but use it wisely**
  - **We've found the AG beacon has too many hosts for effective troubleshooting in the LAN, but it's valuable for load testing**
  - **Consider joining the ESnet beacon on thorn.es.net (much smaller; DOE sites only)**
- **Develop a close working relationship with your service provider**
  - **We're fortunate in this respect; Mike O'Connor and Joe Burrescia have been outstanding**
- **Identify 1 or 2 multicast experts at your vendor's TAC, and open all multicast-related cases with them**
  - **Accept no substitutes!**

Getting ahead of the curve

- Ideally we'd like to be much more proactive when it comes to multicast troubleshooting
- The dream: a multicast "early warning system"
  - Periodically polling routers, performing basic sanity checks on buffers, multicast routing tables, MSDP caches
  - Testing for the specific failure modes we've repeatedly encountered (inductive approach, grounded in operational experience)
- Funding is unclear, but we're currently implementing an "MSDP alarm" as proof of concept
  - Compare MSDP caches on two distant routers

A bigger question; conclusion

- **What's to blame for the poor state of multicast code? (Send me your thoughts!)**
  - **Inherent complexity of the protocols?**
  - **Sloppy development & poor QA?**
  - **Lack of critical customer mass?**
  - **Lack of business case, on the vendor's part, for improving quality?**
- **Conclusion: We should try to influence these factors however we can**
  - **By contributing to IETF efforts**
  - **By holding vendors accountable**
- **In the mean time, we need to focus on**
  - **Mitigation (eliminate failure modes)**
  - **Cooperation (share strategies and information)**

Thanks!

- **Questions / comments**
  - **send to GRBell@lbl.gov**